

# Decision Making with Multiple Goals: Extending the Threshold Approach

Richard Washington<sup>1</sup> and John R. Clarke<sup>2</sup>

<sup>1</sup>Philadelphia VAMC & Dept of Computer and Information Science, U of Pennsylvania

<sup>2</sup>Department of Surgery, Medical College of Pennsylvania—Hahnemann Medical School

*Given a set of conditions that may be present in a patient, this paper presents a method for choosing the optimal combination of treatment actions, taking into account interactions among the available therapeutic and diagnostic procedures. The multiple-threshold approach is an extension of the threshold approach for individual conditions, and it offers a precise way of computing the optimal actions. In addition, tests for an example case show a 11–29% increase in the expected utility of treatments using this method.*

## INTRODUCTION

Given multiple possible conditions requiring treatment and multiple possible therapy actions, a significant problem is to select the best combination of actions to cover all conditions. When the therapeutic and diagnostic actions for the conditions interact, the problem becomes increasingly complex, since one action may play a role in the diagnosis or treatment of multiple problems.

Consider a patient with the possibility of both renal and general abdominal injuries. A CT scan can be used as a diagnostic test for each of the possible conditions, but it may be less than optimal for either condition alone. However, given a reasonable likelihood of both conditions, it may be the preferred option to cover both.

The *threshold approach*<sup>1</sup> was developed to choose the optimal treatment for a single condition. In particular, given the options of not treating a condition, treating it directly, or performing a test to decide whether or not to treat the condition, the method identifies the optimal treatment given the probabilities of the conditions.

This paper presents the *multiple-threshold approach*, a method for choosing the optimal action given multiple conditions and actions, extending the threshold approach to the general case. Given the probabilities of a set of conditions, along with therapeutic and diagnostic procedures, the

multiple-threshold approach produces the optimal treatment for all the conditions, taking into account the diagnostic and therapeutic interactions.

The main benefit of this approach is that it provides a precise way of computing and describing diagnostic and therapeutic interactions. In addition, the ability to reason about multiple conditions improves the expected utility of the treatment. Finally, the results can be used to direct information gathering towards those data that have the most impact on the decision-making process.

## THRESHOLD APPROACH

Given a medical condition with one therapeutic action, the decision is to either treat or not treat the condition. If a diagnostic test is available, the decision also includes the possibility of performing the test and then treating or not based on the test outcome. For a single condition, the optimal action depends on the following information:

- the probability of the condition,  $p$ ,
- the penalty (disutility) of leaving the condition untreated,  $\pi$ ,
- the cost (disutility) of the therapeutic procedure<sup>†</sup>,  $\tau$
- the cost (disutility) of the diagnostic test,  $\delta$
- the sensitivity, or true positive rate ( $TPR$ ), of the diagnostic test
- the specificity, or true negative rate ( $TNR$ ), of the diagnostic test.

Note also that the false positive rate  $FPR = 1 - TNR$ , and the false negative rate  $FNR = 1 - TPR$ .

The utilities of the options can be calculated from this information. The utility of treating is

<sup>†</sup>We assume that the therapeutic procedure has the same cost for a patient with or without the condition. The results can be extended to handle the more general case.<sup>2</sup>

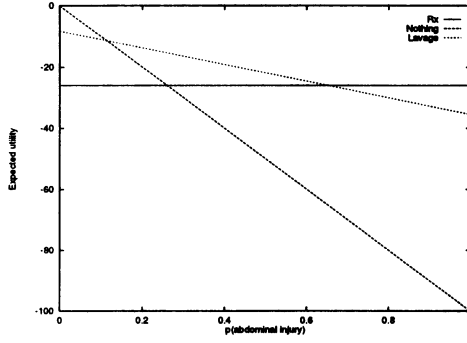


Figure 1: Expected utility given the probability of abdominal injury, from trauma domain.

simply  $\tau$ . The utility of not treating is  $p\pi$ . The utility of acting based on the test outcome is

$$\delta + p \cdot TPR \cdot \tau + p \cdot FNR \cdot \pi \\ + (1 - p) \cdot TNR \cdot 0 + (1 - p) \cdot FPR \cdot \tau.$$

These functions produce three linear equations, as in the example shown in Figure 1. In turn, this describes three regions over which the optimal action is to do nothing, to test, and to treat directly. When the probability of the condition is low, then it is best to do nothing; when the probability is high, then it is best to treat the patient directly; and otherwise it is best to test first.

This basic model can be extended to handle additional therapeutic and diagnostic procedures. In this case there are more utility functions, but again the regions over which one function dominates the others define the probabilities over which one action is optimal.

When there is incomplete information about the true probability of the condition, the probability is only known to lie within an interval. As long as this interval falls within one region, the optimal action remains the same, so no further information gathering is necessary.<sup>3</sup>

## MULTIPLE-THRESHOLD APPROACH

Consider the case where multiple conditions are possible, each with its own probability. There are two sources of interactions among the conditions:

- A single diagnostic procedure may be used to test for multiple conditions. For instance, a CT Scan can identify an abdominal injury or a renal injury.
- A single therapeutic action may be shared among therapeutic procedures for multiple

conditions. For instance, a laparotomy is a component of the therapeutic procedures for treating an abdominal injury or a renal injury.

If a therapeutic or diagnostic action is shared among multiple conditions, the disutility of the overall care is reduced, since the disutility of the action is assessed only once. This may cause a procedure that was considered too morbid for individual conditions to become the best procedure.

For non-interacting conditions, conditions that share neither therapeutic nor diagnostic actions, the utility of actions is found by simply summing the utilities of actions for each of the conditions. In this case, the optimal actions are not affected by the presence of multiple conditions.

When conditions interact, the optimal actions are found by computing the utilities and probabilities of combinations of actions. If we denote by  $u(a)$  the utility of a single action  $a$ , where a therapeutic procedure is a set of actions, then we can define the combined utility of a set of procedures as  $U(P)$ , where

$$U(P) = \sum_{a \in \bigcup_{p \in P} p} u(a), \quad (1)$$

in other words the sum of the utilities of the actions after eliminating duplicates. The combined probability is determined by multiplying the probabilities of the test results (true positive, true negative, false positive, false negative) together.

For each condition, the possible actions are grouped into the three possible types: treat, do nothing, or test. There may be multiple therapy procedures, and there may be multiple diagnostic tests. For each possible combination of interacting actions—one per condition—there is a corresponding formula that describes the expected utility of the combination. The algorithm for determining that formula is shown in Figure 2. Each combination of actions may be further decomposed into the possible test outcomes, since each test may have a positive or negative outcome. So the total utility of the combination of actions is the sum of the utility for each possible set of test outcomes multiplied by the probability of those outcomes.<sup>2</sup> When summed over all outcomes, the utilities of the “do nothing” and therapeutic actions can be removed from the sum and calculated just once, as they are in steps 1 and 2 of Figure 2. The probability of each set of test outcomes is computed in step 3c, and the utility is computed in step 3b, with their product computed in step 3d. The final sum is the sum over all “do nothing” utilities (step

1. Accumulate the “do nothing” actions. Sum the formulas (of the form  $\pi_i \cdot p_i$  for condition  $i$ ).
2. Accumulate the therapeutic procedures. Sum the utilities as in Equation 1.
3. Accumulate the diagnostic tests. For each possible combination of outcomes:
  - (a) Find the optimal therapy procedure for each condition with a positive test outcome.
  - (b) Sum the therapy utilities, eliminating duplicates and all therapy actions already used in step 2. Add to this the penalty  $\pi_i$  for each test  $i$  with a false-negative outcome.
  - (c) Multiply formulas for outcome probabilities (of the form  $a_i p_i + b_i$  for condition  $i$ ) together.
  - (d) Multiply the outcome probability from step 3c by the outcome utility from step 3b.
4. Sum all formulas computed in step 3d to determine the total utility of the test actions.
5. Add the formulas from steps 1, 2, and 4 to determine the final formula.

Figure 2: Algorithm for calculating the expected utility of a combination of actions for interacting conditions.

1), all therapeutic actions (step 2), and all sets of test outcomes (step 4).

The algorithm above produces a formula describing the expected utility of a combination of actions, where the conditions interact. Not all conditions will interact, so the full algorithm first groups conditions into *interaction sets*—sets of conditions that interact—then computes formulas for each interaction set and sums them to determine the final set of formulas. If interaction between two conditions is defined as above, namely that the conditions share either a therapy action or a diagnostic procedure, then the interaction set is found by taking the transitive closure of the interaction relation: if a condition  $i$  interacts with a condition  $j$  in the set, then  $i$  is in the set as well.

The full algorithm is the following:

1. Calculate interaction sets  $I_1, I_2, \dots, I_m$  as described above.
2. For each  $I_j$ , compute a set of formulas  $F_j$  as in Figure 2.
3. For each  $\{f_1, \dots, f_m\} \in F_1 \times F_2 \times \dots \times F_m$ , compute the sum:  $f = \sum_{j=1}^m f_j$ .

The formulas  $f$  computed in the final step are the expected utilities of actions.

Given  $n$  conditions, each formula is the combination of terms of up to  $n$  variables, but no variable has an exponent greater than 1. So the resulting formula is of the form:

$$f(p_1, \dots, p_n) = \sum_{i_1=0}^1 \dots \sum_{i_n=0}^1 a_{i_1, \dots, i_n} p_1^{i_1} \dots p_n^{i_n} \quad (2)$$

where each  $i_k$  is 1 when variable  $k$  is present in the term and 0 if variable  $k$  is absent, and where

$a_{i_1, \dots, i_n}$  is an arbitrary coefficient of the term. The formula is thus the sum of a constant term, a set of terms with single variables  $p_j$ , a set of cross terms with 2 variables  $p_j p_k$ , a set of cross terms with 3 variables  $p_j p_k p_l$ , ..., a single cross term with  $n$  variables  $p_1 \dots p_n$ . This has some interesting mathematical properties:

- The projection of the function to any single axis is linear: for any constants  $c_i$ ,  $f(c_1, \dots, c_{k-1}, p_k, c_{k+1}, \dots, c_n) = ap_k + b$ .
- From this it follows that two functions intersect at most once along a single axis. The intersection of two functions may be nonlinear, but the property holds in any case.
- Following from that, if a function  $f_i \geq f_j$  at each of the vertices of the hypercube  $[0, 1]^n$ , then  $f_i \geq f_j$  at all points within the hypercube. In fact this is true for any hypercube in  $\mathbb{R}^n$ .

The last property allows formulas that are provably worse than others to be eliminated efficiently. Given the set of formulas, there can be a large number generated that are provably worse than others over the probability space  $[0, 1]^n$  (the set of all legal probability values for each condition). Empirically we have eliminated 50–60% of the candidate formulas using this property.

Given the set  $\mathcal{F}$  of formulas that represent the expected utilities of the possible actions for each condition, the remaining step is to determine which combination of actions is optimal for a given set of probabilities on the conditions. In fact, this is the step that needs to be computed on-line; the formula computation, which is NP-complete, can be computed ahead of time. The optimal actions for a given set of probabilities  $p_1, \dots, p_n$  are

merely the actions corresponding to the function with maximum utility at that point:

$$f_{opt} = \arg \max_{f \in \mathcal{F}} f(p_1, \dots, p_n).$$

This is the  $n$ -dimensional equivalent of the threshold approach presented earlier: there the goal was to find the maximum function for a single given probability  $p$ ; here it is for a set of probabilities.

## EXPERIMENTAL RESULTS

To test the method, we extracted a subset of the knowledge base of TraumAID,<sup>4</sup> a system for assistance in trauma management. See Figure 3 for the data used. The utilities range from  $-100$  to  $0$ , and the disutility scale was evoked from four demographically different surgeons using the standard gamble technique. Note that conditions 1–3 interact, and conditions 4–5 interact.

We considered three test cases. Case I uses conditions 1 and 2. Case II uses conditions 1–3. Case III uses conditions 1–5. For Case I, the formulas are shown in Figure 4(a) and the optimal actions are shown in Figure 4(b). Note that the diagnostic action “CT Scan” appears prominently in the combined actions, although it would not for either condition individually. This is a result of the sharing of the test costs. Table 1 contains the results for all of the cases. The overall utility of a method was determined by numerical integration over the  $[0, 1]^n$  space. The two methods considered were the multiple-threshold approach and the threshold approach applied to each of the conditions individually. As can be seen from the table, the more interaction present in the case, the greater the gain from the multiple-threshold approach.

## DISCUSSION

The *multiple-threshold approach* has been presented as a method for precisely determining the optimal actions to take given multiple interacting conditions. Tests for an example case show an 11–29% increase in the expected utility of treatment for this approach over an approach that ignores interactions.

Given the complex multi-dimensional and non-linear nature of the formulas, the explicit computation of the exact thresholds remains an outstanding issue, requiring multi-dimensional geo-

metric representation and reasoning. We have not needed to compute the thresholds explicitly, since the relevant computations are to determine the optimal actions at particular points within the probability space. The implicit representation of the threshold suffices for the computations we need and foresee.

There may be some “bedside” information that is easy to determine (negligible cost), but may change the probability of a condition.<sup>3</sup> As long as the range of possible probabilities falls within one region (as mentioned for the simple threshold approach), the treatment will remain the same, so no further information gathering is necessary. If the range crosses a threshold, then this can be used to determine which information would be most useful to gather. In the case of multiple conditions, this will direct the information-gathering process to those conditions whose probability needs to be more precise before the treatment can be localized within one region. The coordination of bedside information gathering with the multiple-threshold approach will be the next enhancement of the method.

## Acknowledgements

The first author was supported by a Veterans Affairs Medical Informatics Fellowship, administered through the Philadelphia VAMC.

## References

1. S. G. Pauker and J. P. Kassirer. The threshold approach to clinical decision making. *New England Journal of Medicine*, 302:1109–1117, 1980.
2. R. Washington and J. R. Clarke. Determining thresholds for interacting decision strategies. In preparation, 1996.
3. J. R. Clarke and Jr. O'Donnell, T. F. A scientific approach to surgical reasoning. vi. thresholds and confidence. *Theoretical Surgery*, 6:177–183, 1991.
4. B. Webber, R. Rymon, and J. R. Clarke. Flexible support for trauma management through goal-directed reasoning and planning. *Artificial Intelligence in Medicine*, 4(2):145–163, April 1992.

1. Abdominal Injury. Penalty for missed injury = -100.  
Rx: Rx nonspecific [Laparotomy (-25) + Colostomy (-1)]  
Dx: Peritoneal Lavage, cost -7, sensitivity 0.97, specificity 0.95  
Dx: CT Scan Abdomen, cost -8, sensitivity 0.98, specificity 0.99
2. Renal Injury. Penalty for missed injury = -42.  
Rx: Observe [Observe (-9)]  
Rx: Inspect [Laparotomy (-25) + Inspect (-5) + Urologist consultation (-1)]  
Dx: IVP, cost -9, sensitivity 0.99, specificity 0.99  
Dx: CT Scan Abdomen, cost -8, sensitivity 0.98, specificity 0.99
3. Duodenal Injury. Penalty for missed injury = -41.  
Rx: Repair [Check allergies (0) + Antibiotics (-1) + Laparotomy (-25) + Duodenum repair (-15)]  
Dx: CT Scan Abdomen, cost -8, sensitivity 0.98, specificity 0.99
4. Tracheal Injury. Penalty for missed injury = -100.  
Rx: Perform trachea repair [Perform thoracotomy (-30) + Trachea repair (-15)]  
Dx: Bronchoscopy, cost -15, sensitivity 0.98, specificity 0.98
5. Bronchial Injury. Penalty for missed injury = -100.  
Rx: Perform bronchus repair [Perform thoracotomy (-30) + Bronchus repair (-15)]  
Dx: Bronchoscopy, cost -15, sensitivity 0.98, specificity 0.98

Figure 3: Data used for empirical evaluation of multiple threshold approach. Case I uses data items 1–2, Case II uses items 1–3, and Case III uses items 1–5. Utilities range from -100 to 0.

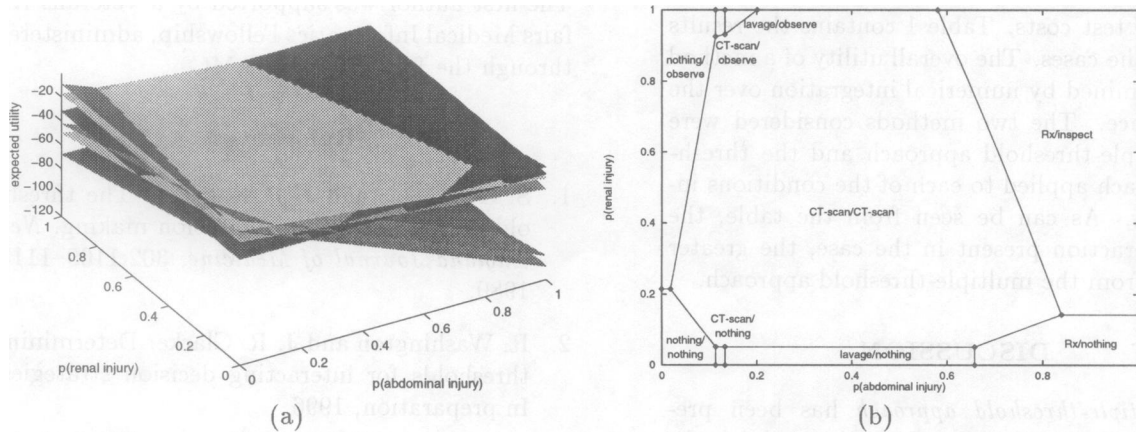


Figure 4: Combinations of actions for Case I of the test data. (a) Surfaces produced by formulas. (b) Optimal actions given probability of injuries. The regions at the top and bottom of (b) have been expanded in the figure for visibility.

Table 1: Empirical results of the threshold and multiple-threshold approaches for the test cases, comparing expected utilities of the multiple-threshold and threshold approaches.

Case	conditions	interaction sets	multiple-threshold utility	threshold utility	improvement
I	2	{1,2}	-24.62	-27.71	11%
II	3	{1,2,3}	-39.18	-55.20	29%
III	5	{1,2,3} {4,5}	-90.37	-121.82	26%